


E96

Introduction to IMS V9 Integrated Online Reorganization

Vern L Watts vlwatts@us.ibm.com



IMS

technical conference

Las Vegas, NV

September 15 - September 18, 2003

© IBM Corporation 2003

User Requirements

- Availability of data should not be affected by reorganization
 - ▶ 100% of data available 100% of the time



Solution Highlights

- OLR is supported for HALDB PHDAM and PHIDAM partitions
- OLR is controlled with IMSplex commands on TSO SPOC or Control Center, or with Classic IMS commands
 - ▶ INIT, TERM, UPD, QRY
- System ITASK process per partition
- Dual data sets during cursor based reorganization
- Concurrent SHAREPLEX updates are allowed while OLR is active
 - ▶ V8 OLR Coexistence SPE will be available to allow full datasharing between V8 and V9 while OLR is active



© IBM Corporation 2003

IMS Technical Conference

- ▶ Online reorganization (OLR) is essentially phase II of HALDB. OLR of a HALDB PHDAM or PHIDAM partition will, upon command initiation, run as a system ITASK in the DLISAS address space.
- ▶ OLR can be activated by a new IMSplex command via the Operations Manager (OM) Application Programming Interface (API) or the Control Center, or by a classic IMS command.
- ▶ OLR is a cursor based reorganization on a paired set of data sets. A cursor marks the point at the end of the database records already copied. As the copying proceeds, this cursor moves forward through the database.
- ▶ Concurrent IMS SHAREPLEX updates are allowed while small amounts of data are moved and reorganized. IMS determines whether to expect a given database record to be in either the input or the output data sets based on whether the record is either beyond or prior to the cursor. This is completely transparent to an application program accessing the database.
- ▶ An OLR Coexistence SPE will be provided on V8 to allow full data sharing between V8 and V9 with OLR active, for HALDBs that are marked ONLINE REORG CAPABLE=YES via DBRC INIT.DB or CHANGE.DB commands.

Solution Highlights ...

- OLR throttling is supported via OLR rate
- Extension of HALDB naming convention
 - ▶ A-J & X will become M-V & Y and back
- Output data sets can be preallocated by user or automatically created by IMS
- No data definitional changes via OLR
- PHIDAM primary index is also reorganized
- All data movement during OLR are logged



- ▶ How fast an OLR runs is determined by the RATE specified on the INITIATE OLREORG or UPDATE OLREORG commands, system resources and system utilization. The amount of updates to the IMS log and system contention can also affect the speed at which an OLR runs. The RATE can be changed when adjustments are desired depending on system contention (slow down the OLR) or lack of system contention (speed up the OLR).
- ▶ OLR has extended the data definition and data set naming convention established for HALDB. Multiple data set groups in a HALDB database use the characters A-J in the DDNAMEs and data set names of the supported ten data set groups, and the primary index for a PHIDAM database uses the character X in these names. This has been expanded by implementing the characters M-V and Y for the paired set of data sets.
- ▶ Output data sets can be preallocated by the user for size and placement. Otherwise they are automatically created by IMS to match the input data sets.
- ▶ The OLR process does not support data definitional changes. It only provides reclustering and space distribution advantages.
- ▶ OLR for a PHIDAM partition includes reorganization of the corresponding primary index data set also.
- ▶ For possible recovery of either an input or an output data set, all database changes are logged, including the database records copied from the input to the output data sets.

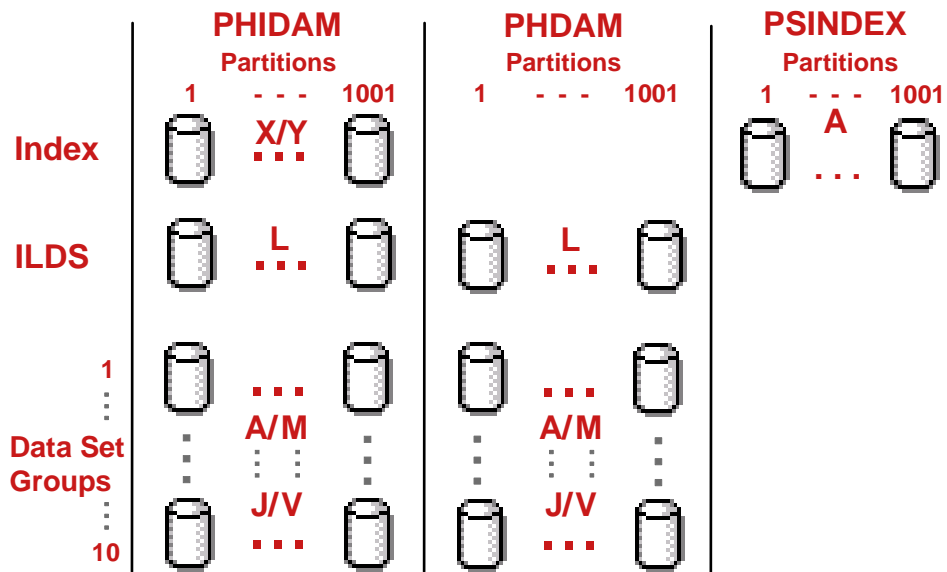
Solution Highlights ...

- Being cursor based
 - ▶ Once the cursor-active status is recorded in RECON, OLR must complete or an offline reorg must be run to reset the OLR settings in RECON
 - Offline Reorg reload always loads the A-J side
 - ▶ HALDB partition cannot be accessed unless both the input and output sets of data sets are available
- OLR input data sets and their contents become obsolete at reorg completion
 - ▶ Image copy and recovery records no longer applicable to "dead" data



- ▶ Before the OLR starts, there is a single active set of data sets for the HALDB partition - either the A-J and X set or the M-V and Y set. The data sets in the other set contain no useful information. Once the OLR initialization process has completed, the active set of data sets is treated as the input set and the other (paired) set becomes the output set. At the end of the OLR initialization process, the OLR of the partition is recorded in the RECON with a cursor-active status.
- ▶ Once the cursor-active status has been recorded, the partition is comprised of both the A-J and X set of data sets as well as the M-V and Y set. At this point the OLR must complete or an offline reorganization must be run in order to reset the online reorganization settings. (Note that offline reorganization reload always load the A-J side.) Also, the partition cannot be accessed unless both sets of data sets are available.
- ▶ After all database records have been copied, the output data sets become the active set and the input data sets become the inactive set. Unless an initial load or an offline reorganization is done, successive OLRs for the partition then alternate between the two sets of data sets.
- ▶ At the completion of OLR, the input data sets and their contents become obsolete. Image copy and recovery records are no longer applicable to "dead" data. All prior image copies and log data of "dead" data sets are rendered void.

Naming Convention



IMSplex Command - INITIATE OLREORG

- INIT OLREORG NAME(partname1,partname2)
SET(RATE(100)) OPTION(DEL)
 - ▶ Initiate or resume OLR



- ▶ INITIATE OLREORG causes OLR to begin or resume on the partition or partitions specified. RATE(100) is the default which means run the OLR as fast as possible with no coded waits. OPTION(DEL) is the default and indicates the deletion of the old input data sets is to be attempted as part of the termination processing when the OLR completes. Note that one big advantage of specifying the NODEL option is to gain better control of data set placement.
- ▶ When the OLR completes, a message indicating this will be sent to the system console. The command output for this command is defined in XML and is available to automation programs which communicate with OM.

IMSplex Command - TERMINATE OLREORG

- TERM OLREORG NAME(partname1,partname2)
 - ▶ Temporarily stop OLR



© IBM Corporation 2003

IMS Technical Conference

- ▶ TERMINATE OLREORG causes OLR to be stopped for the specified HALDB partition(s). OLR will no longer be active or have an owning IMS. OLR can be resumed on any IMS, with the INITIATE OLREORG command.
- ▶ OPTION(FORCE) indicates that the OLR for the named partition is to be stopped when the next record boundary is encountered. All moved data up to that point is committed to DASD and no backout is required. OPTION(ABORT) causes the OLR to be stopped immediately. Backout may be required depending on the state of the OLR.
- ▶ If an OLR is stopped prior to completion, the NODEL keyword is **not** retained and will need to be specified on the INITIATE OLREORG command that is issued to resume the stopped OLR or on the UPDATE OLREORG command.
- ▶ The command output for this command is defined in XML and is available to automation programs which communicate with OM.

IMSplex Command - UPDATE OLREORG

- UPD OLREORG NAME(partname1,partname2) SET(RATE(50))
OPTION(NODEL)
 - ▶ Change the rate and/or delete options



- ▶ UPDATE OLREORG allows the changing of the RATE and/or DELETE options for one or more OLRs.
- ▶ The command output for this command is defined in XML and is available to automation programs which communicate with OM.

IMSplex Command - QUERY OLREORG

- QRY OLREORG NAME(partname1,partname2)
STATUS(RUNNING) SHOW(ALL)
 - ▶ Return status information about active OLRs



© IBM Corporation 2003

IMS Technical Conference

- ▶ QUERY OLREORG will return the status of one or more OLR tasks depending on the specified keywords. If NAME(partition name) is not specified, then the default NAME(*) will be used and the responses will only be for HALDB partitions that have OLR in progress.
- ▶ STATUS() selects which OLRs to return status for. STATUS(RUNNING) specifies that the output will be for the OLRs that are owned by each IMS (to which the command is routed) for the specified partition name(s). SHOW() specifies the output fields to be returned. SHOW(ALL) returns the rate at which the OLR is running and also the **number of bytes moved** to the output data set.
- ▶ The command output for this command is defined in XML and is available to automation programs which communicate with OM.

IMSplex Command - QUERY DB

- QRY DB STATUS(OLR)
 - ▶ Shows all partitions with OLR in progress



- ▶ The QUERY DB command is enhanced to show OLR status. QUERY DB STATUS(OLR) will return all partitions with OLR in progress. No rate information is returned. This is returned in the QUERY OLREORG command output.

Classic Commands for OLR

- All IMSplex OLR commands (except QUERY) are supported as classic commands
- Command syntax for /INITIATE, /TERMINATE, /UPDATE the same as for IMSplex commands
 - ▶ Only one partition name can be specified
 - ▶ Commands passed to IMS AOI exit and logged to secondary master
- Command response in message format sent to inputting terminal
 - ▶ INITIATE OLREORG
 - DFS058I INITIATE COMMAND IN PROGRESS followed by DFS0725I message to system console
 - ▶ TERMINATE OLREORG and UPDATE OLREORG
 - DFS0725I message



© IBM Corporation 2003

IMS Technical Conference

- ▶ The commands to initiate, terminate and update OLR are supported as classic commands. /INITIATE OLREORG, /TERMINATE OLREORG and /UPDATE OLREORG may be issued from terminals, AOI applications, EMCS consoles and APPC/OTMA programs. They are also passed to the AOI user exits and logged to the secondary master. Note that /QUERY OLREORG is not supported.
- ▶ The classic command syntax is the same as the syntax of the equivalent IMSplex command from OM API. The classic commands are however preceded by a CRC, usually a '/. Note that only one partition name can be specified on the classic commands.
- ▶ When a command is issued from OM API, the command response is encapsulated in XML tags. When the command is entered as a classic command, the command response is returned in the message format.
- ▶ INITIATE command - If successful results in the DFS058I INITIATE COMMAND IN PROGRESS message followed by the DFS0725I message to the system console.
- ▶ TERMINATE and UPDATE command - result in the DFS0725I message to the inputting terminal.
- ▶ Any errors result in the following error messages:
 - DFS0725I (new) - INITIATE|UPDATE|TERMINATE OLREORG COMMAND FOR DB dbnamexx COMPLETE. CC= nn where dbnamexx is the HALDB partition name and nn is the completion code.
 - DFS107 - REQUIRED KEYWORD NOT PRESENT
 - DFS110 - COMMAND KEYWORD OLREORG INVALID FOR yyyyyyy where yyyyyyy is TRACKER or XRF ALTERNATE
 - DFS128 - TOO FEW KEYWORDS
 - DFS136 - COMMAND xxxxxxxxxxxx INVALID FOR yyyyyyy where yyyyyyy is TRACKER or XRF ALTERNATE
 - DFS165 - COMMAND VERB IS INVALID
 - DFS2026 - No WKAP STORAGE
 - DFS2038 - INSUFFICIENT STORAGE AVAILABLE TO PROCESS COMMAND
 - DFS2262 (new) - TOO MANY DATABASE PARAMETERS
 - DFS3630 - GENERIC PARAMETER NOT ALLOWED

Classic Commands for OLR ...

- /DIS DB OLR
 - ▶ Shows all partitions that have OLR in progress locally
 - ▶ This is the classic command flavor of QUERY OLREORG
 - ▶ Command response in new format mapped by DISPLAY Dsects D18 and D77
- /DIS DB
 - ▶ Enhanced to show OLR status for specified partitions that have OLR in progress locally



- ▶ /DIS DB OLR command can be considered as the classic command flavor of QUERY OLREORG. /DIS DB OLR returns all partitions that have OLR in progress locally. The following information is returned: master DB name, partition name, OLR rate, number of bytes moved and OLR status.
- ▶ Response is in a new format and is mapped by DISPLAY format ids D18 and D77.
- ▶ The attribute OLR cannot be specified with any other attribute such as STOPPED or ALLOCS on the /DIS DB OLR command. If specified, the command is rejected with a new DFS135I message.
- ▶ If a /DIS DB OLR command is issued on the RSR tracker or the XRF alternate, the command is rejected with the DFS132 message.
- ▶ The /DIS DB command is enhanced to show the OLR status for the specified partitions or databases which have OLR in progress locally.

Classic Commands for OLR ...

```
■ /INIT OLREORG NAME (PDHDOKA) SET (RATE (5))
DFS058I 12:08:07 INITIATE COMMAND IN PROGRESS
DFS2970I - OLR STARTED FOR NAME=PDHDOKA
DFS0725I INIT          OLREORG COMMAND FOR DB PDHDOKA  COMPLETE. CC=    0

■ /DIS DB OLR
DATABASE PART      RATE      BYTES  STATUS
DBHDOK01 PDHDOKA    5        86776  RUNNING
*03099/120858*

■ /DIS DB PDHDOKA DBHDOK01
DATABASE TYPE  TOTAL UNUSED  TOTAL UNUSED  ACC  CONDITIONS
DBHDOK01 PHDAM                UP
PDHDOKA PART                UP  ALLOCS, OLR
DBHDOK01 PHDAM                UP
PDHDOKA PART                UP  ALLOCS, OLR
PDHDOKB PART                UP  NOTOPEN
PDHDOKC PART                UP  NOTOPEN
PDHDOKD PART                UP  NOTOPEN
*03056/140011*
```



Output Data Set Creation

- Use data set if it exists
- Space in OSAM blocks or VSAM records equivalent to input data set
- SMS-managed
 - ▶ Same storage class as input data set
 - ▶ Same number of volumes as input data set
 - ▶ With guaranteed space attribute, primary space allocation taken on all volumes
- Non-SMS, non-VSAM
 - ▶ UNIT=SYSALLDA - data set created on storage volume or public volume
 - ▶ Multi-volume data sets not created automatically
- Non-SMS, VSAM
 - ▶ Same volume(s) as input data set




© IBM Corporation 2003

IMS Technical Conference

- ▶ To simplify the OLR process for a HALDB partition, each of the output data sets can be created **automatically** by IMS. For each data set group defined in the DBD and for the primary index of a PHIDAM database, the output data set is created if it doesn't already exist as a cataloged data set. The indirect list data set is not created automatically because there is not a corresponding output version of it.
- ▶ In order to reserve approximately the same amount of space that was reserved for the input data set regardless of the DASD types involved, the output data set's space is requested as a number of OSAM blocks or VSAM records.
- ▶ If an input data set is SMS-managed, then the corresponding output data set is SMS-managed as well, and the storage class used to create the output data set is that of the input data set. If the input data set has extents on only one DASD volume, the output data set is also created on a single volume. If the input data set has extents on more than one DASD volume, the output data set is created to allow the same number of volumes. If the storage class of the input data set has the guaranteed space attribute, the primary space allocation will be taken on each of the volumes when the output data set is created.
- ▶ If an input data set is non SMS-managed and is non-VSAM, the output data set is created as though UNIT=SYSALLDA has been specified on a DD statement. This causes the data set to be created on a storage volume or, if no storage volume is available, on a public volume. Note that if the input data set is a multi-volume data set, the output data set will not be created automatically.
- ▶ If an input data set is non SMS-managed and is VSAM, the output data set is created on the same volume or volumes as the corresponding input data set.

Input Data Set Deletion

- At OLR completion
 - ▶ Old input data sets deleted (all A-J and X data sets or all M-V and Y data sets)
 - ▶ Deletion unless OPTION(NODEL) is specified on INIT or UPD command



- ▶ As part of the successful completion of an OLR, all of the input data sets involved in the reorganization become inactive. These data sets may be deleted automatically by IMS, regardless of whether they were created by IMS or were pre-existing data sets. This automatic data set deletion can be suppressed by specifying the NODEL option on the INIT OLREORG or UPD OLREORG command.

DBRC

- DB record for HALDB master contains
 - ▶ ONLINE REORG CAPABLE=YES|NO
- DB record for HALDB partition contains
 - ▶ OLR active indicator
 - ▶ SSID of IMS owning OLR
 - ▶ A-J or M-V indicator for active set of data sets
 - ▶ ONLINE REORG CAPABLE=YES|NO



- ▶ The DB record for a HALDB master database contains new information for OLR - a field indicating whether OLR is enabled for the HALDB master and its partitions.
- ▶ The DB record for each HALDB partition contains new information for OLR - a field indicating whether OLR is processing, the SSID of the IMS currently performing an OLR, a field indicating which set of DBDSs is active for the partition and if the M-V data sets exist, and a field indicating whether OLR is enabled for the partition.

DBRC ...

- DBDS records in RECON for both A-J and M-V data sets
 - ▶ A-J and M-V data sets are treated as one logical set of DBDSs
 - ▶ Recovery information in the 2 DBDS headers are maintained as one logical DBDS
- REORG record
 - ▶ Contains information for online or offline reorganization
 - ▶ Recorded under both A-J and M-V DBDSs
 - ▶ Has OLR start time
 - ▶ Has OLR stop time
 - ▶ Has RECOV=YES|NO



- ▶ The M-V DBDS records are defined in RECON with the exact same attributes of the corresponding A-J DBDS records. The A-J and M-V data sets are treated as one logical set of DBDSs. The recovery information in the two DBDS headers are maintained as one logical DBDS. For example, GENMAX and CAGRP are the same. Any changes done to the active DBDS are also done to the inactive DBDS.
- ▶ The REORG record for a HALDB partition DBDS contains new information for OLR. For an offline reorganization, the database must be /DBR'd. Therefore only the start time of the reorganization is necessary. With OLR, the database is not /DBR'd, so the duration of the process is important. The stop time of the reorganization is now recorded in the REORG record as well as a new indicator showing that this is an OLR and whether the OLR may be used as input to recovery. A REORG record is recorded under both the A-J and the M-V DBDSs. The REORG record for the DBDS that was the output data set will indicate it can be used as input for recovery. The REORG record for the DBDS that was the input data set will not indicate it can be used for recovery. It is recorded to indicate to DBRC that reorganization occurred and full recovery for this data set is not allowed after OLR has completed.

Utilities

- Image Copy
 - ▶ Copies A-J or M-V data sets
 - No need for user to be aware of which set of data sets is active
 - ▶ Not allowed while OLR is active
 - ▶ Coding of DD statement is **not** required
 - GENJCL no longer generates DD statement
 - Data sets are dynamically allocated



- ▶ Image copy will copy the HALDB partition A-J or M-V data sets. The new function will be transparent to the user. The user does not need to be aware of which set of data sets is the active set and does not have to provide this information. Image Copy will determine if the M-V data sets are to be copied instead of the A-J data sets.
- ▶ While OLR is active for a partition, none of the three DB Image Copy utilities can copy any of the data sets for that partition. Even if the TERM OLREORG command is used to temporarily stop OLR, an image copy still cannot be made for any data set in that partition.
- ▶ Since dynamic allocation is recommended, GENJCL no longer generates DD statements for the HALDB DBDS being copied. It is no longer necessary to code a DD statement in the JCL when copying partition data sets as they will be dynamically allocated.

Utilities ...

- Recovery
 - ▶ Recovers A-J as well as M-V data sets
 - Timestamp recovery to middle of OLR is not allowed if OLR completed
 - ▶ Log-only recovery for output data sets
 - Does not require prior image copy
 - GENJCL.RECOV prepares recovery from start of OLR



- ▶ DB Recovery Utility DFSURDB0 will recover HALDB partition A-J data sets as well as M-V data sets.
- ▶ Timestamp recovery (TSR) to a point where a HALDB OLR was active is not allowed if OLR has completed.
- ▶ Once the cursor-active status for OLR is recorded in RECON, recovery of any of the output data sets with DFSURDB0 can be accomplished solely from DB change records (type x'50' log records) until OLR has completed and a subsequent image copy has been made. An empty output data set is considered to be the starting point from which DB change records can be applied.
- ▶ The RECON reflects the beginning of OLR as a starting point from which forward recovery of the output data sets can be done, even after OLR is finished. Until an image copy has been made for the output data set, the DBRC GENJCL.RECOV prepares recovery of an output data set from this point even though no physical image copy exists.

Utilities ...

- Change Accumulation
 - ▶ Accumulates changes for A-J as well as M-V data sets
 - ▶ DBRC GENJCL.CA command treats start of OLR as the most recent image copy and accumulate changes to the output data set from that point
- Batch Backout
 - ▶ Recognize OLR PSB by its name
 - Numeric character '0' followed by the 7-byte partition name
 - Must be run in DLI (not DBB) region to build OLR PSB



© IBM Corporation 2003

IMS Technical Conference

- ▶ Change Accumulation Utility DFSUCUM0 will accumulate changes for HALDB partition A-J as well as M-V data sets.
- ▶ The DBRC GENJCL.CA command treats the start of OLR as though it were the most recent image copy and causes changes to the output data set to be accumulated from that point.
- ▶ A PSB is dynamically built and used for OLR. The PSB does not exist in PSBLIB or ACBLIB. The name of the PSB is constructed by prefixing the 7-byte partition name with a numeric character '0'. When it is necessary to run Batch Backout Utility (DFSBB000) to backout inflight OLR work, DFSBB000 will recognize the named PSB is from OLR because it begins with a numeric character which is not allowed for defined PSBs. Note that DFSBB000 must be run in a DLI region in order for the OLR PSB to be built.

Restart and XRF

- Shutdown - all active OLRs are suspended before a shutdown checkpoint and resumed on next restart
- Abnormal termination, XRF takeover - inflight OLRs dynamically backed out to last sync point and OLRs resumed on restart/takeover



- ▶ All active OLR tasks are suspended before a shutdown checkpoint. They will be resumed on the next restart after the OLR environment is re-initialized.
- ▶ All OLR tasks that were active when IMS abnormally terminated will have the inflight OLR(s) dynamically backed out to the last sync point. When backout is complete, the OLR environment will be rebuilt and OLR tasks will be resumed at restart. The same applies to XRF takeover in which case the OLR tasks will be resumed on the new active IMS.

RSR

- No need for DB to be taken offline when OLR is recognized by the tracking IMS
- Tracking site RECON updated with OLR related information
 - ▶ New counters to keep track of outstanding OLRs
 - ▶ Do not reflect OLR ownership
- Create or reuse output data sets
 - ▶ If database level tracking and partition is defined as DBTRACK
- No automatic OLR resumption at takeover /ERE



© IBM Corporation 2003

IMS Technical Conference

- ▶ Remote Site Recovery (RSR) supports OLR. Unlike offline reorganization, there is no need for the databases to be brought offline when OLR is recognized by the RSR tracking IMS.
- ▶ As part of normal RSR tracking, the tracking site RECON is updated with OLR related information in both Recovery Level Tracking (RLT) and Database Level Tracking (DLT) mode. If the tracking IMS is running with the DLT feature and the partition is defined as DBTRACK in RECON, the output data sets are automatically reused if they exist. If they do not exist, they will be created appropriately as either A-J or M-V set of data sets. The tracking site RECON will also reflect the current status of covered databases. New counters keep track of outstanding OLRs that have been recorded in RECON but not yet applied to the databases.
- ▶ If data set characteristics at the active site are manually changed for partitions undergoing OLR, then they must also be manually changed at the tracking site.
- ▶ OLRs that were active at the time of RSR takeover will not be automatically resumed at the takeover /ERE. After the ILDS and primary index data sets are being rebuilt by DFSPREC0, INITIATE OLR command needs to be issued to resume the OLRs.

Log Records and Logging

- Existing record changes
 - ▶ X'50' database change log record
 - X'80' bit in DCB number to identify M-V and Y data sets
 - Odd RBAs in M-V data sets
- New records introduced
 - ▶ X'29' log record type for OLR in DFSLOG29
 - 2910 - OLR owned
 - 2920 - UPD OLREORG command
 - 2930 - Output data set creation info
 - 2940 - Cursor active
 - 2950 - OLR unit of reorg commit (cursor movement)



© IBM Corporation 2003

IMS Technical Conference

- ▶ The x'50' database change log record has been changed. The DCB number field has x'80' bit on for M-V data sets. Odd RBAs are used in M-V data sets while even RBAs are used in A-J data sets. This means a 4GB data set size limit for HALDB partitions.
- ▶ x'20'/ x'21' (DB open/close) log records and x'27' (DB extension) log records have also been changed. The DCB number field has x'80' bit on for M-V data sets.
- ▶ New x'29' log record type is introduced to reflect the progress of OLR processing.

Log Records and Logging ...

- Statistics for OLR are in the x'2950' log records
 - ▶ Segments moved in this UOR
 - ▶ Bytes moved in this UOR
 - ▶ Total segments in the partition moved before this UOR
 - ▶ Total bytes in the partition moved before this UOR
 - ▶ Roots moved in this UOR
 - ▶ Locks held by this UOR
 - ▶ Start time of this UOR
 - ▶ Execution time of this UOR
 - ▶ Time interval waited before this UOR
- Insertion into new output data sets is logged as x'50' database change log records
 - ▶ Log volume during OLR will increase



© IBM Corporation 2003

IMS Technical Conference

- ▶ OLR statistics counters are included in the x'2950' record which is logged at the end of each unit of reorganization (UOR). The statistics include :
 - Number of segments moved in this UOR
 - Number of bytes moved in this UOR
 - Total number of segments in the partition moved before this UOR
 - Total number of bytes in the partition moved before this UOR
 - Number of roots moved in this UOR
 - Number of locks held by this UOR
 - Start time of this UOR
 - Execution time of this UOR
 - Time interval waited before this UOR
- ▶ The logging volume attributable to OLR is important to understand because OLR of a partition generates x'50' log records for insertion into the new output data sets. OLR of multiple partitions in parallel could generate sufficient log data to impact normal transaction processing. The OLR RATE can be changed when adjustments are desired depending on system contention.

Performance Considerations

- OLR throttling is supported via OLR rate
- If shareplex environment, dedicate IMS to run OLR
 - ▶ Buffer pool definition can be tuned for OLR
 - No flushing of data from other workloads
 - ▶ IMS logs dedicated to OLR
 - Log volume will not impact current workload
 - OLR log data will not be included in recovery and CA of non-OLR data



- ▶ The speed of OLR is determined by the RATE specified on the INITIATE OLREORG or UPDATE OLREORG commands, system resources and system utilization. The amount of updates to the IMS log and system contention can also affect the speed at which an OLR runs. The RATE can be changed when adjustments are desired depending on system contention (slow down the OLR) or lack of system contention (speed up the OLR).
- ▶ In a Shareplex environment (SHRLEVEL 2 or 3), it is beneficial to dedicate a separate IMS to run OLR for performance reasons.

Migration Considerations

- V7/V8 to V9
 - ▶ CHANGE.RECON UPGRADE
 - Pre-existing HALDBs will be marked ONLINE REORG CAPABLE = NO
 - ▶ New HALDBs defined on V9 are OLR capable by default (ONLINE REORG CAPABLE = YES)
 - ▶ New partitions to HALDB will inherit ONLINE REORG CAPABLE setting from the HALDB master
 - ▶ CHANGE.DB OLRCAP or INIT.DB OLRCAP to enable OLR for HALDB master on V9
 - ▶ CHANGE.DB OLRNOCAP or INIT.DB OLRNOCAP to disable OLR for HALDB master on V9



© IBM Corporation 2003

IMS Technical Conference

- ▶ Migration from V7 or V8 to V9 --
 - When the CHANGE.RECON UPGRADE command is issued, all existing HALDBs are marked ONLINE REORG CAPABLE=NO.
 - New HALDBs defined on V9 are by default OLR-capable. This is true for definitions via either the PDU or Batch commands. The setting takes effect on the HALDB master database, meaning that all partitions in that HALDB will be OLR-capable.
 - New partitions added to a HALDB will inherit the OLR-capable setting for the HALDB master database.
 - To enable the OLR-capable setting for a HALDB master and its partitions, the CHANGE.DB OLRCAP or INIT.DB OLRCAP command should be used. Note that for V9, OLRCAP is the default setting.
 - To disable the OLR-capable setting for a HALDB master and its partitions, the CHANGE.DB OLRNOCAP or INIT.DB OLRNOCAP command should be used.

Coexistence Considerations

- V7 and V9
 - ▶ OLR is allowed on V9 for all HALDBs that are marked ONLINE REORG CAPABLE = YES
 - ▶ V7 will **not** see HALDBs that are marked ONLINE REORG CAPABLE = YES



- ▶ Coexistence between V7 and V9 --
 - OLR is allowed to be run on V9 for all HALDBs that are marked OLR-capable.
 - V7 will not be able to see any HALDB that is marked OLR-capable.

Coexistence Considerations ...

- V8 and V9
 - ▶ OLR is allowed on V9 for all HALDBs that are marked ONLINE REORG CAPABLE = YES
 - ▶ With OLR Coexistence SPE applied
 - V8 will see all HALDBs (including HALDBs that are marked ONLINE REORG CAPABLE = YES)
 - V8 will be able to share data with V9 during OLR
 - Once data sharing begins, must use V9 utilities except for Archive, Batch Backout and Log Recovery
 - ▶ Without OLR Coexistence SPE applied
 - V8 will **not** see HALDBs that are marked ONLINE REORG CAPABLE = YES



© IBM Corporation 2003

IMS Technical Conference

- ▶ Coexistence between V8 and V9 --
 - OLR is allowed to be run on V9 for all HALDBs that are marked OLR-capable.
 - With the OLR Coexistence SPE applied, V8 will be able to see all HALDBs (including HALDBs that are marked OLR-capable). V8 will also be able to share data with V9 while OLR is active. Once data sharing has begun between V8 and V9, the V9 utilities must be used except for Archive, Batch Backout and Log Recovery.
 - Without the OLR Coexistence SPE applied, V8 will not be able to see any HALDB that is marked OLR-capable.